

Application No. 10/763,135
Amendment Dated October 24, 2011
Reply to Office Action of May 9, 2011

AMENDMENTS TO THE CLAIMS

The below listing of claims will replace all prior versions and listings of claims in the application:

Listing of Claims:

1. - 34. (Canceled)

35. (Currently amended) The method of claim 64, wherein a first application workload executes on a first server cluster having a first domain and the remote domain includes a second domain having a second server cluster running a second application workload further comprising:

monitoring execution of the first application workload to determine whether a the performance requirement[[s]] for execution of the first application workload specified in the service level agreement will continue to be met; and

responsive to a determination that the performance requirement[[s]] for execution of the first application workload will not continue to be met, sending a request to the second domain to assign one or more of the plurality of server nodes in the second server cluster at the second domain to the execution of the first application workload.

36.-37. (Cancelled)

38. (Previously presented) The method of claim 64, further comprising:
receiving a refusal from the remote domain denying the request, the refusal from the remote domain being based upon evaluation of an impact on a service level agreement negotiated for an application workload executing on the remote domain if the request was granted.

39. (Previously presented) The method of claim 35, wherein responsive to a determination that the performance requirements for execution of the first application workload are

Application No. 10/763,135
Amendment Dated October 24, 2011
Reply to Office Action of May 9, 2011

not being met, the method further comprises:

sending a request to a third domain to assign one or more of a plurality of server nodes in a third server cluster at the third domain to the execution of the first application workload, wherein a third application workload is executing on the third server cluster at the third domain while the first application workload is executing on the first server cluster at the first domain, and

wherein the request sent to the third domain specifies a number of server nodes requested, a duration in which the number of server nodes requested will be needed, and a monetary value associated with the request.

40. (Previously presented) The method of claim 39, further comprising:

receiving a first counter offer from the second domain, the first counter offer specifying one or more of a different number of server nodes than the number of server nodes requested, a different duration for the number of server nodes requested, and a different monetary value for the request; and

receiving a second counter offer from the third domain, the second counter offer specifying one or more of a different number of server nodes than the number of server nodes requested, a different duration for the number of server nodes requested, and a different monetary value for the request, wherein what is specified in the first counter offer differs from what is specified the second counter offer.

41. (Previously presented) The method of claim 40, further comprising:

sending a first response to the second domain, the first response indicating acceptance or rejection of the first counter offer from the second domain; and

sending a second response to the third domain, the second response indicating acceptance or rejection of the second counter offer from the third domain.

42. (Previously presented) The method of claim 39, further comprising:
receiving a refusal from the third domain denying the request, the refusal from the third domain being based upon evaluation of an impact on a service level agreement negotiated for the third application workload executing on the third server cluster at the third domain if the request was granted.

43. (Cancelled)

44. (Previously presented) The method of claim 64, wherein the monetary value associated with the request is a payment amount for the number of server nodes requested.

45. (Previously presented) The method of claim 64, wherein the first application workload is a transaction application workload.

46. (Previously presented) The method of claim 45, wherein the transaction application workload comprises stock trades.

47. (Previously presented) The method of claim 64, wherein the remote domain executes a parallel application workload.

48. (Previously presented) The method of claim 47, wherein the parallel application workload involves optimization of a stock portfolio.

49. (Previously presented) The method of claim 64, wherein the threshold performance requirement comprises a throughput requirement.

50. (Previously presented) The method of claim 64, wherein the threshold performance requirement comprises a response time requirement.

51. (Previously presented) The method of claim 64, wherein the threshold performance requirement comprises an availability requirement.

52. (Currently amended) The method of claim 64, wherein the threshold performance requirement ~~level agreement~~ comprises a downtime requirement.

53. (Cancelled)

54. (Currently amended) The method of claim 64, wherein monitoring execution of the application workload ~~on the subset of server nodes assigned to execute the first application workload is monitored, comprising~~ comprises:

monitoring one or more of a transaction rate, a transaction response time, availability of a server node, and utilization of a server node.

55.-57. (Cancelled)

58. (Previously presented) The method of claim 64, further comprising receiving an acceptance of the request in accordance with a comparison of a monetary value and a minimum acceptable payment amount.

59. (Previously presented) The method of claim 58, wherein the minimum acceptable payment amount is determined in accordance with a determination of the value of processing operations performed at the remote domain.

60. (Previously presented) The method of claim 59, wherein the value of processing operations performed at the remote domain is determined in accordance with a service level agreement of the remote domain.

61. (Currently amended) A method for determining whether to support an application workload at a local cluster of nodes using a resource at a remote domain remote from the local cluster, the method comprising:

receiving at the remote domain from the local cluster a request for at least one server node determined at the local cluster in accordance with a threshold of performance requirements ~~of the local cluster wherein the request specifies a number of nodes requested, a time duration for which the requested nodes are needed, and a monetary value determined in accordance with a monetary penalty amount specified by a service level agreement for a failure to meet a requirement of the service level agreement;~~

transmitting from the remote domain to the local cluster an acceptance of the request in accordance with the monetary value; and

allocating the at least one server node in accordance with a the monetary penalty amount for a the failure to meet the threshold performance requirement specified by of the service level agreement.

62. (Currently amended) An article of manufacture, comprising:

at least one non-transitory computer usable medium having computer readable program code embodied therein for performing a method for supporting an application workload using a resource at a remote domain, the computer readable program code means in said article of manufacture comprising:

computer readable program code means for causing a computer to assign a subset of a plurality of server nodes to execute the application workload;

computer readable program code means for causing a computer to execute the application workload on the assigned subset of the plurality of server nodes;

computer readable program code means for causing a computer to monitor execution of the application workload to determine whether a threshold of a performance requirement of a service level agreement specifying performance requirements and a monetary penalty amount for execution of the application workload is met;

computer readable program code means for causing a computer to send a request for at least one server node to the remote domain wherein the request specifies a number of nodes requested, a time duration for which the requested nodes are needed, and a monetary value associated with the request responsive to a determination that the threshold of the performance requirements is not being met wherein the monetary value is determined in accordance with the monetary penalty amount specified by a service level agreement for a failure to meet a requirement of the service level agreement;

computer readable program code means for causing a computer to receive from the remote domain an acceptance of the request in accordance with the monetary value; and

computer readable program code means for allocating the at least one server node in accordance with the monetary penalty amount for the failure to meet the requirement of the service level agreement.

63. (Cancelled)

64. (Currently amended) A method for supporting an application workload using a resource at a remote domain, the method comprising:

monitoring execution of the application workload to determine whether a threshold performance requirement of a service level agreement is met wherein a monetary penalty amount for not meeting the threshold performance requirement is specified by a service level agreement;

responsive to a determination that the threshold performance requirement of the service level agreement is not met, sending a request for a resource;

receiving a counter offer from the remote domain specifying a resource differing from the requested resource; and

allocating the resource specified by the remote domain in the counter offer according to the monetary penalty amount ~~service level agreement~~.

65. (Previously presented) The method of claim 64, wherein the request for the resource includes a request for at least one server node and the counter offer specifies a different number of server nodes than the number of server nodes requested.

66. (Previously presented) The method of claim 64, wherein the request for the resource includes a request for a time duration for which the resource is requested and the counter offer specifies a different time duration.

67. (Currently amended) The method of claim 64, wherein the request for the resource includes a monetary value determined in accordance with the monetary penalty amount ~~for a failure to meet a requirement of the service level agreement~~ and the counter offer specifies a different monetary value.

68. (Previously presented) The method of claim 64, wherein the steps of monitoring execution of the application workload, sending the request for the resource, receiving the counter offer and allocating the resource specified in the counter offer are performed by an on-demand manager at a local domain executing the application workload.